

EXHIBIT 7
[FILED UNDER SEAL]

Message

From: [REDACTED]
Sent: 5/18/2014 2:03:28 AM
To: [REDACTED]
Subject: Re: Another crazy idea to boost AdX revenue

Hey [REDACTED],

thanks a lot for your sympathy. I badly need a break, I'm going on vacation this coming week. How about we set something up for the week of May 26?

[REDACTED]

On Sat, May 17, 2014 at 2:46 PM, [REDACTED] > wrote:

[REDACTED],

I would like to talk to you about your rant :-)

I fully sympathize with it. Let's set up a meeting.

[REDACTED]

On May 17, 2014 10:28 AM, "[REDACTED]" > wrote:

Hey [REDACTED],

It is a vexing problem, isn't it? We've known for a long time that sum of winning bids on AdX is almost twice the sum of transaction prices. The question is, how to capture some of this surplus and transfer it over to the seller.

Selective information hiding to increase competition thickness is one way to increase revenue. Another is running a lottery designed to threaten the highest bidder you'll drop them from the auction with some probability unless they bid above a "premium service" reserve price. See my writeup, and link to a paper from msft describing a randomized auciton mechanism they call BIN-TAC.

<https://docs.google.com/a/google.com/document/d/1ihM9rGBL5zzNiJ8NkGUcFB6gxJDmjo9cK5WdnLgrL2A/edit?usp=sharing>

You can study Myerson's theory on revenue optimal auction design. The issue is that observed bid distributions don't satisfy the usual regularity / monotone hazard rate conditions, so optimal Myerson mechanism is necessarily randomized. Bid distributions are also not identical, so you want to have a way to discriminate among the bidders based on your knowledge of their valuations / preferences.

If you allow yourself to personalize the mechanism to buyers, I claim the easiest change to implement and launch is personalizing reserve prices. Based on my toy model [REDACTED], and the model is really toy -- something I hacked together in a week. The rest is getting lawyers, managers, PM's and engineers (who have been previously told by PM's that automated reserve pricing is out of question) on the same page, and trying to keep concerns from folks like [REDACTED] about adwords cpd drop in perspective. Not much effort has gone into actual engineering of the model.

<Warning: Rancid Rant follows. Stop reading if you're weak on the stomach.>

Unfortunately we have much more basic problems here in NYC than worrying about auction design. In a culture obsessed with producing code and launches, It is very difficult to get people rewarded for producing and sharing ideas, producing good designs, and saying no to bad ones. Many experienced people have left, and of those who are still around are looking for an exit. As a replacement we're getting a stream of nooglers

who have no clue why things work the way they do, and need to be taught what a mapreduce is (let alone what is an auction) before they can be of much use.

FYI, [REDACTED] [REDACTED].
Purely as an example, I am really close to running an experiment with dynamic RPO. The last nice to have before running it is a TNT test. I could block off two days writing it and stop doing more important things like making sure people on my team are leading rewarding and productive lives. I could also ask another engineer who is much better at this sort of thing than I am. I am however scared shitless to ask him for help, because then I'll be accused of not producing enough code, and if he writes the CL, there will be less documentation of my contribution. There are only so many notches on the perf scale and I'd really not enjoy being put on a performance improvement plan. It is difficult to work on meaningful projects in an environment like that.
<end of rant>

On May 17, 2014 7:16 AM, "[REDACTED] [REDACTED]" wrote:

I can't sleep...

Some data to show that there is some room for this to work...

I pulled all filled AdX publisher requests, where the auction_discount > 80%... and grouped by buyer_network_id (AdWords = 1).

[REDACTED]

[REDACTED]


[REDACTED]


Here is the first many networks... in case anyone is curious:

[REDACTED]

34543









On Fri, May 16, 2014 at 9:29 AM, [REDACTED] > wrote:

[REDACTED]. We'll sync with you in a separate thread once we understand more...

On May 16, 2014 9:25 AM, [REDACTED] wrote:

[REDACTED]: can you report the magnitude of the bids you see (the winning bid vs the losing adwords bid boosted by 25%)?

On Fri, May 16, 2014 at 12:20 PM, [REDACTED] wrote:

Independent of how the AdX buyer computed their bid, it's surprising that quite often GDN doesn't have an ad worth [REDACTED]. Gtrade team is looking into this further.

On May 16, 2014 7:17 AM, [REDACTED] wrote:

Alok, regarding your experiment where RTB values impressions higher. Could this be caused by buyers essentially using the same bidding logic across multiple exchanges, where most others have no real second price auction? [REDACTED]

On May 16, 2014 5:42 AM [REDACTED] > wrote:

[REDACTED] So yes, we do see some crazy bids in the wild. In general, most queries have very little spread between winning bid and transaction price, but a few

[REDACTED] looked into doing pricing per cookie, where we could set a reserve price based on past bidding behavior on that cookie. [REDACTED]. The downside of his approach though is that one needs to build and maintain a table to store pricing information per cookie, which is fairly expensive both in terms of engineering labor and machine cost. Per buyer pricing is just much easier to implement.

On Thu, May 15, 2014 at 9:51 PM [REDACTED] wrote:

REDACTED - PRIVILEGE

On Thu, May 15, 2014 at 2:20 PM, [REDACTED] wrote:

The continuum is:

"show thousands of cookies and tell the buyer that you get one of them"..... "show two cookies, one is true", "show the cookie"

The beginning of this continuum is the world before RTB. Adnetworks were negotiating the average CPM with publishers for each segment of the traffic. Networks were assessing user base (cookies) using pixeling or other available statistics. Back in early days of RTB discussions, before Google enters into this space, it was argued that RTB removes market inefficiency by valuing each individual impression based on the unique features of that impression. This might be the right solution (I have to think more), but it would be ironic to go backward in the continuum. May be we should go back all the way!

On Thu, May 15, 2014 at 9:54 AM, [REDACTED] > wrote:

Is there is a continuum of options between "show the cookie" and "show two cookies, one is true" and "show nothing." [REDACTED]

On Thu, May 15, 2014 at 9:14 AM, [REDACTED] > wrote:

+1, this is sure to be a fruitful effort regardless...

On Thu, May 15, 2014 at 8:50 AM, [REDACTED] > wrote:

FYI gdn is running a short experiment [REDACTED]. Even then AdX buyers win a lot of queries. It would be interesting to know why an AdX buyer values the query [REDACTED] what GDN thinks, given the large inventory of keyword/contextual/not cookie ads that Google can show. I recall [REDACTED] was doing some investigation here (any update?) If we understand why GDN has low demand on some queries it may shed light on the wider density issues.

On May 15, 2014 8:43 AM, [REDACTED] > wrote:

Agree with the bundling. I think we could definitely pull it off and I'm willing to suffer the slings and arrows from buyers [REDACTED]

=)

On 15 May 2014 11:39, [REDACTED] > wrote:

Drew I don't think the plan would be to say we "don't know". It would be to say "we're not telling you because it's good for the pub". If we package it with something that's good for buyers [REDACTED] as you suggested) it might make sense.

At this point we all agree the proposal isn't yet ready to withstand the brutality of the real world, so maybe let it develop a bit first.

On Thu, May 15, 2014 at 11:36 AM, [REDACTED] > wrote:

So bad choice of words on my part. It isnt lying but buyers would want to know why on earth cookie cant figure out what the real cookie is.

[REDACTED]

On 15 May 2014 11:31, [REDACTED] wrote:
Sending no cookie information has a reduced impact on their RTB systems, but I agree that "I'm going to give you two scenarios and flip a coin about which one arises" isn't lying.

On Thu, May 15, 2014 at 8:29 AM, [REDACTED] > wrote:
No, more precisely... my goal is to increase bid density in the auction, in an incentive compatible way that still keeps efficiency very close to maximum.

[REDACTED]: I'm not sure why sending an RPC with the semantics of "this request has a 50% chance of being user A and 50% chance of being user B" is interpret able as lying. We're going to try hard to make it true (e.g. half the time they get this query, it will be A and have the time it will be B). We'll be sending them (effectively) 2x as many queries to make this happen, of course...

How is this any more/less deceptive than sending them no cookie information? It seems equivalent to me...

On Thu, May 15, 2014 at 8:19 AM, [REDACTED] > wrote:
The biggest issue we have here is that we would knowingly be sending a fake cookie to buyers that would break any targeting or profiling they are doing. If it was an option to also send similar audiences as a feature, great, but instead were deliberately misleading retargeters on who the user is.

[REDACTED]

On 15 May 2014 11:12, [REDACTED] > wrote:
[REDACTED] if your goal is to increase revenue at the expense of a CPD hit (i.e. keep system value constant, but induce a higher cash transfer from buyers to sellers), dynamic reserve price optimization in my opinion is a much more straightforward proposition than playing information games with the buyers. In an ideal world we should do both, but given that we have a limited supply of engineers you should help me launch dynamic RPO first.

[REDACTED]

P.S. And of course if we're going for impact, we should not forget to get dynamic revshare out of the door.

On Thu, May 15, 2014 at 11:02 AM, [REDACTED] wrote:

[REDACTED]

The impact to the buyers should be neutral convs/bid, and volume should stay about the same (across all advertisers). They will definitely take a CPD hit though, which should pretty much entirely be a function of higher auction competition.

[REDACTED]

On May 15, 2014 7:54 AM, [REDACTED] wrote:
+adx-quality, [REDACTED]

[REDACTED]: thank you, this is a fruitful direction.

As Preston says hiding information in the BidRequest is much less controversial than intentionally sending fabricated information. Publishers already do it in the form of bid requests for anonymous inventory. I think we could totally get away with running experiments on this -- the first field we thought of hiding/modifying is the reserve price field.

[REDACTED]. Thanks to a bunch of work by [REDACTED] we see the logjam gradually breaking up. We'll soon be able to ask for explicit publisher permission to run experiments, and I'm hoping one day our dear eng/pm leadership decides it's OK to run a few cross-exchange experiments without undue risk of alienating our customers and regulators.

If someone wants to propose what specific BidRequest fields to hide, from what bidders and under what conditions please talk to us -- we'd want to chat about some kind of offline estimate of impact and eng work required to run and analyze a live experiment.

[REDACTED]

On Thu, May 15, 2014 at 10:07 AM, [REDACTED] wrote:
There is a simpler way than [REDACTED] two cookie auction. If we can forecast the likely competitiveness of the auction fairly well, we do the following:

If pretty competitive ([REDACTED]), we reveal the cookies and let the auction proceed normally.

If concentrated ([REDACTED]), we conceal the cookie from the bidders. These impressions go into a pool of "minimal information opportunities". We might provide a little information about them, e.g. geo only, provided that doesn't concentrate many submarkets.

This way we don't have to hold two auctions, and explain to the bidders that we are running hypothetical auctions.

None of the theory has been tested empirically; this is a pure theory exercise. Details are here:

<http://vita.mcafee.cc/PDF/targeting.pdf>

On Thu, May 15, 2014 at 6:34 AM, [REDACTED] > wrote:
[REDACTED]

PRIVILEGED AND CONFIDENTIAL

REDACTED - PRIVILEGE

On Thu, May 15, 2014 at 9:26 AM, [REDACTED] wrote:

One more thought: if you are going to send fabricated data to buyers, you'll need to send different fake cookies to different buyers. If you do that suddenly you have to be very careful lest you open yourself to charges of unfair/unequal treatment. Lawyers tend to be very good at not understanding brilliant arguments about how the math shows you are in fact being fair if it is in their interest not to. I'm sure its a solvable problem given enough resources, but I'm not sure the combined resources of people on this list are sufficient to solve it.

[REDACTED]

[REDACTED]

If you received this communication by mistake, please don't forward it to anyone else (it may contain confidential or privileged information), please erase all copies of it, including all attachments, and please let the sender know it went to the wrong person. Thanks.

AKAP

[REDACTED]

--

[REDACTED]

--

[REDACTED]